



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Comparing Imaging and Non-Imaging Techniques for Reducing Background Clutter and Resolving Distant Point Sources

R. Wurtz, K.-P. Ziock, L. Fabris, R. Graham

November 11, 2005

2005 IEEE Nuclear Science Symposium and Medical Imaging Conference
Fajardo, PR, United States
October 23, 2005 through October 29, 2005

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Comparing Imaging and Non-Imaging Techniques for Reducing Background Clutter and Resolving Distant Point Sources

Ron Wurtz, Klaus-Peter Ziock, Lorenzo Fabris, *Member, IEEE*, and Rion Graham

Abstract— To reach maximum sensitivity, any method used to search for orphan sources must be insensitive to local variations of the background. Using imaging and non-imaging techniques, we analyzed the same data acquired by a search instrument deployed as a large-area, coded-mask imager. Data from many passes past a 1 mCi source at 65 m from the instrument were used to construct a model of the instrument response. We then used the model to “hide” the source in data taken in a light urban environment. We compared the success of detecting the hidden sources using imaging coded-mask methods, pseudo-imaging based on a zero-area matched filter, and non-imaging using simple thresholding. The results clearly indicate the superiority of imaging with the coded-mask techniques returning the best results.

I. INTRODUCTION

There is currently a great deal of interest in the detection of orphan sources [1, 2]. Such sources can be considered as any radiation source that is not where it belongs, either by accident or malicious intent. One approach to finding orphan sources is to sweep a suspect region for the source with radiation detectors. Such a wide area search for nuclear materials preferably finds all such materials and makes no false alarms. Unfortunately, the background fluctuations in a search are usually so large that one must either use a high threshold to avoid false alarms, thereby missing faint sources, or use a low threshold and suffer many false alarms. It is obvious that the raw rate of events in a detector makes a poor distinction between sources and background variations [3, 4]. Using coded-aperture masks, the Large Area Imager (LAI) [4] encodes the data coming into a mobile detector so that background fluctuations can be suppressed and the signature of point sources, even those at a distance, can be properly flagged.

Manuscript received November 11, 2005. This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

R. Wurtz is with Lawrence Livermore National Laboratory, Livermore, CA 94550 USA (telephone: 925-423-8504, e-mail: wurtz@llnl.gov).

K.-P. Ziock is with Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA (telephone: 865-574-0272, e-mail: ziockk@ornl.gov).

L. Fabris is with Lawrence Livermore National Laboratory, Livermore, CA 94550 USA (telephone: 925-423-1258, e-mail: fabrisl@llnl.gov).

R. Graham is with Lawrence Livermore National Laboratory, Livermore, CA 94550 USA (telephone: 925-424-3580, e-mail: graham48@llnl.gov).

Two questions remain: first, are there methods of processing non-imaged data that reveal point sources against the fluctuations, and second, if such methods exist, how does their sensitivity compare with imaging? We have made a direct comparison between imaging and non-imaging techniques by analyzing the same data from the LAI, first as an imager and then as a non-imager using both a simple threshold and a zero-area matched filter.

II. INSTRUMENT

The Livermore Large Area Imager is a 3 by 19 array of 10 cm by 10 cm by 10 cm NaI(Tl) detectors located behind a coded-aperture mask. It is described in detail in [4]. For these measurements the system was mounted in the back of a truck (Fig. 1) and was configured to observe the curb side of the street. It uses custom electronics to acquire list-mode data from each of the detectors while the system is in motion. Included in the data stream are “events” from a Doppler-radar-based fifth-wheel that indicate the linear location of the imager from a fixed starting point. Timestamp events are also included at one millisecond intervals to allow one to correct for dwell time at a given location.

The instrument generates one-dimensional images of the radiation field observed at any given location. Based on data from the fifth-wheel, the images are mapped to overall world maps that can be used to find the location of sources. Parallax effects are used to generate world-maps for each of ten different 10-meter deep distance bins set on the central 5-meter ranges from the imager (i.e. 5 m, 15 m, 25 m, etc.)



Fig. 1. The Large Area Imager mounted in the back of a rental truck. The mask (A) projects a shadow onto the pixilated detector (B). From the shadow and the known mask pattern, 1-d images are combined with the truck location and parallax effects to generate 2-d images as shown in Fig. 2.

The resulting 2-d plot (Fig. 2) reveals the location of point sources while nulling large-scale background variations.

The data are normalized to the dwell time at a given location and counting statistics errors are propagated through the analysis routine as explained in [4]. The data is also analyzed as the counts per meter obtained without using the imaging analysis. This data is the same as would be obtained with a source of half the intensity if a normal (no mask) detector were used since the mask blocks 50% of the radiation from the source. The results of the imaging analysis centered on 65 meters are compared with the counts per meter data analyzed using pseudo-imaging and a threshold approach as described below.

III. METHODOLOGY

A. Real and Simulated Data

Even if one has enough counts in a detector from an orphan source to make a statistical detection, one must still deal with background variations as a function of location. Thus, to determine the efficacy of a detection system, it is insufficient to make many trials at one location and determine if the source is found. Rather, one should hide a source in many locations and see how many times one finds it, misses it, and gets false alarms while looking for it; however, restrictions on using sources in public make this approach difficult.

To overcome this problem, we took high fidelity data on a source at LLNL, generated the system response, and then used this to obtain a “model-source” that could be located anywhere in no-source data sets taken in the world-at-large. For the model source, data were acquired by placing a 1 mCi ^{137}Cs source 65 m from a road and driving the LAI past it at 16 km/h. A total of 20 data sets were collected, 10 with the source present and 10 without.

To better reject artifacts in coded aperture images, each data set was acquired in two integrations, one with a mask and one with an anti-mask configuration. To swap between the two configurations, the open and closed elements of the mask are interchanged. The LAI was designed to make this interchange easy. (A new instrument currently under construction will use two inline imagers to obtain the two data sets in a single pass.) In addition to removing unwanted artifacts from the images, the net effect of using two passes is that it doubles the effective detector area of the imager from 0.57 m^2 to 1.14 m^2 .

For the analysis, only data around the 662 keV peak of the ^{137}Cs source is used. The detectors are all calibrated to the same energy scale, and an energy window from 590 to 730 keV was selected. For the world-at-large data, a total of about 16 minutes of data were collected in the Livermore area at the same velocity and with the same energy cuts. Both a mask and an anti-mask pass were used as well. The model source was then hidden at different locations in these data.




Fig. 2. The imager “world-view” 2-d image for half of the source data used to generate the model. The source is detected at nearly 11 sigma. Subsequent “image” data uses only the image results centered at 65 m from the detector.

B. Receiver Operating Characteristic Curves

We use Receiver Operating Characteristic (ROC) curves to determine performance. A ROC curve plots the probability of detection as a function of the probability of a false alarm. The two probabilities are generated by varying the detection threshold and counting the number of detections at source locations and the number of false alarms at non-source locations. The two parameters are then plotted against each other, eliminating the common threshold. A poor detector gives a diagonal line from (0,0) to (1,1). A perfect detector gives a right angle plot going from (0,0) to (0,1) and then to (1,1), i.e. all detections are made without a single false alarm.

IV. IMAGING MODE

A. Imaged Source Model

The imaged source model is taken from the LAI world-image at 65 m (Fig. 2.) To include possible effects of imaging in the model, the image itself is not used. Rather, the model is generated and it is saved before the imaging step in the analysis. This means that we save the detector map and associated counting statistics data for each of the pixels in the image at 65 m. The summed, no-source, detector maps are subtracted from the summed, source, detector maps at each location before an image conversion is normally made. This difference is saved to disk.

To add the model to a data set, the model detector maps are scaled for source strength (0.1 for a 1 mCi source) and then added to the data at each location of the data set for a




Fig. 3. Imaging model used in the analysis. The blue line represents the uncertainty in the model at each location.

selected source location before it is imaged. The source adds to the data for about 75 m on either side of the central location. The imager field of view spreads this another factor of two, as shown in Fig. 3.

B. Model Performance

The imager source model is quite good, we compared an actual source run with a no-source run that has the model added, and found the model has slightly higher counting statistics noise since we do not use a “noiseless” fit. Also, no adjustment has been made to shift the starting points of different runs to optimize the overlap. Finally, the effect of trees in the real data (see below) remains in the model, making the detection weaker than for an unobstructed source.

C. Imager Detection and False Alarms

Ultimately we wish to determine the probabilities of detection and false alarm as a function of pixel height above zero as the threshold, computed in units of standard deviation (sigma), is varied for the imaged data. Curves of these items as a function of the confidence level (sigma) are shown in Fig. 4.

The probability of false alarm is generated from a total of 644 no-source pixels examined in the world-at-large data. This represents 16.8 minutes of driving time at 16 km/h.

To obtain the probability of detection, the appropriately scaled model source is added to the data from the world-at-large. To keep the model source locations statistically independent, the model is inserted into these data every 75 m for a total of 42 locations. Results from these pixels are used to generate the probability of detection. The curve for a 1.0 mCi source is shown in Fig. 4. Values of 0.5 mCi and 1.5 mCi were also run.

D. Imager ROC Curves

ROC curves for the three model source strengths generated for the imager are shown in Fig. 5. Note that the same probability of false alarm data is used for each of the different strength sources to generate the ROC curves. The performance at 1.5 mCi is perfect, indicating that more trials




Fig. 4. Probability of false alarm and detection as a function of the sigma value for a pixel.




Fig. 5. ROC curves for the fully imaged data. The data for the 1.5 mCi source is “perfect” indicating that no false alarms are obtained before all of the sources are detected.

are needed to properly map the response. The 1.0 mCi result is only slightly worse.

V. NON-IMAGING MODE

A. Distant Source Flux Fitting Function and Matched Filter

The intensity from a point source to a detector at distance d and detector inclination angle θ , goes as

$$I(d; \lambda, \theta) \propto \frac{1}{d^2} e^{-d/\lambda} \cos \theta$$

where λ is the mean free path in the atmosphere.

The perceived intensity for the detector at location x on a straight path having a distance of closest approach b , past a source at x_0 becomes

$$I(x; I_0, x_0, b, \lambda) = I_0 \frac{b}{d^{3/2}} e^{-d/\lambda}$$

where

$$d \equiv \sqrt{(x - x_0)^2 + b^2}.$$

To account for this shape factor and its variation with b in analyzing non-imaged data, we selected the zero area matched filter (ZAMF) method. A ZAMF is constructed by making a kernel the exact shape of the item you are looking for and then subtracting its total area. The subtraction gives it the zero-area property. This means that, when convolved with the data, the kernel adds no net counts, reducing the contribution from all spatial variations that do not match the model. The ZAMF is then convolved with the data and the result gives a measure of the intensity of any source it matches at each point along the intensity path [5].

If no sources are present, the values of the convolution should be normally distributed. If a source that matches the filter is present, then the value of the convolution at the




Fig. 6. Source + background minus background, and model fit and residuals.

point corresponding to the center of the source should exceed the expected fluctuations in the convolution if the source is bright enough. This gives us a distance-specific method to optimize the search for sources in non-imaged data. Since the ZAMF varies with distance to the source, it can be used to construct 2-d images similar to that shown in Fig. 2 from the full image analysis. For this reason this approach is considered a pseudo-imaging technique.

B. Non-Imaging Model Extraction

Fig. 6 shows the non-imaged data from the LAI (filled and unfilled diamonds). The data is the sum of all 10 source data runs minus the sum of the 10 no-source data runs. The source is located approximately at -70 m along the path.

The peak shows a few of the issues encountered with search data. First, the data are quite noisy. Second, despite our “ideal” source placement, it turns out that there were two one-meter thick trees within 5 meters of the source. These shielded the source from the road and are evident as notches in the peak. To remove these effects from the model, we made a 5-parameter fit that takes into account the distance of closest approach, the geometric detector area, and atmospheric attenuation, as well as the peak height, center and baseline offset. We exclude the data obviously behind the trees (the unfilled dots) to obtain a fit (solid




Fig. 7. One-fifth of the world-at-large data in raw counts. The red line is the same data with the model source added at zero meters along the track.




Fig. 8. Filtered world-at-large data. The red line is the same data with the model source added and then filtered.

line) that was in good agreement with the known parameters, i.e. the location was ± 3 m from measured, the distance of closest approach was ± 9 m from truth, and the atmospheric attenuation length was 131 ± 49 m, which compares favorably with the expected 105 m.

C. Sample Non-Imaged Data

We acquired about five kilometers of no-source background (world-at-large) data. The typical count rate in our detector between 590 and 730 keV was 200 counts per meter at 16 km./h. The “real” fluctuations along a path, which correspond to background emission, shielding, and geometry, are about $\pm 20\%$. In Fig. 7, the black line is one section of such data. The red line is the same data with the noiseless model inserted at the location of zero meters for a 1 mCi source. This peak is hidden both by the noise and the local fluctuations and would be impossible to find without *a priori* information on the background.

D. Sample Matched Filter Processed Data

Fig. 8 shows the result of applying the ZAMF for a source at 65 m to the raw data in Fig. 7. The ZAMF is based on the theoretical response expected for a source at 65 m, not from the noiseless fit used to make the model. This means the ZAMF is slightly different from the inserted shape. Nevertheless, a much larger excursion is seen for the plot with the inserted source (red) than for that with no source (black).

E. Non-Imager Detection and False Alarms

Because of local variations, the raw data are not distributed normally, and so it is difficult to establish a useful value of confidence level for alarming. However, for ROC analysis purposes, the distribution of counts without hidden sources must be compared to the distribution with hidden sources. Further, we would like to see that the matched filter ameliorates this unfortunate sensitivity to local variations and more closely yields a normal distribution so that standard statistical techniques may be used.

Fig. 9 shows the distribution of “detections” above a threshold for world-at-large data with and without the model peak. The left figure is for the raw data, the right figure is for matched-filter-processed data. Both plots are




Fig. 9. Fraction of counts above threshold as a function of threshold. On the left are curves for raw counts with and without source, on the right are the same curves for filtered data

for 1 mCi sources. Notice how a low threshold means many false alarms, particularly in the raw data, while a high threshold reduces the false alarm rate, but also misses more sources.

F. ROC Curves For Raw Data and Processed Data

Now we can compute the ROC curves for the non-imaged data. The reported source strengths are compensated for the 50% reduction in signal due to the mask. This enables us to make a quantitative comparison for the value of imaging based on the important parameter of how much time and energy are expended chasing alarms as compared with the success rate.

Fig. 10 shows the ROC curves for the non-imaged data. The “raw counts” ROC (left) shows that at an acceptable detection rate—say, above 99%—even for a 1.5 mCi source we would get a false alarm indication 4 out of every 10 meters.

The ZAMF-processed data is much better. It suggests that a non-imager can be used to detect and locate a few mCi point source at several tens of meters. However, at 65 m we still have a 15% false alarm rate for 99% detection of a 1 mCi source.

VI. DISCUSSION

A practical interpretation of the ROC curves is that they describe how often one must suffer false alarms in order to successfully detect an orphan source of a given intensity at a given distance. Because we only have 5 km of world-at-large data, we are limited in our knowledge of false alarm rates that happen fewer than once every 20 min. at 16 km/h. This would be unacceptable in scenarios where each false alarm requires 20 minutes of follow-up but might be acceptable if each alarm took only one minute. Thus, to fully interpret the implications of a given ROC curve, the operational scenario must be well defined. Also, since real detections are rare, it would clearly help to have more data so one could better map the upper left corner of the ROC curves.

For the non-imaging modes, the false alarm rate is computed as the fraction of total distance traveled above threshold. The computation is more properly performed by determining the number of times per distance traveled that the




Fig. 10. On the left, ROC curves for raw count thresholding for three different sources. On the right, ROC curves for ZAMF thresholding when the source distance matches the filter.

measured intensity value crosses the threshold. The imaging ROC curve more correctly takes care of this problem because its response is mapped to a single pixel. For the non-imaging methods, we plan to apply a sequential probability ratio test to derive ROC curves that can be more directly compared with the imaging result.

This analysis was limited to the 1-d map at a fixed 65-m distance from the road. For both the imaging mode and the ZAMF, a 2-d map may be computed, displaying the sigma value above background for a number of different ranges. However, the parameter returned by the ZAMF that we present here is not normalized to allow comparison of true sigma among the computed ranges.

VII. CONCLUSION

The results indicate that the zero area matched filter improves search performance for distant point sources beyond that possible with normal, non-imaged data. However, the same detector used in mask/anti-mask imaging mode reduced the false alarm rate by a factor of at least ten for 99% detection of a 1.5 mCi source at 65 meters. The superiority of the imager is unmistakable despite the “imperfect” point source model used.

REFERENCES

- [1] T. Hjerpe, C. Samuelsson, “A comparison between gross and net count methods when searching for orphan radioactive sources,” *Health Physics*, 84, 203, 2003.
- [2] D.L. Upp, R.M. Keyser, “Performance of a car-mounted neutron and gamma-ray monitoring system for illicit material detection,” INMM 45th Annual meeting. Available from: http://www.ortec-online.com/papers/cartop_inmm04.pdf.
- [3] K.P. Zioc, W.H. Goldstein, “The lost source, varying backgrounds and why bigger may not be better,” in *Unattended Radiation Sensor Systems for Remote Applications*, J. Trombka, D. Spears, P. Solomon ed., AIP Conf. Proc. 632, (American Institute of Physics, Melville, NY, 2002) 60-70.
- [4] K.P. Zioc, W.C. Craig, L. Fabris, R.C. Lanza, S. Gallagher, B.K.P. Horn, N.W. Madden, “Large area imaging detector for long-range, passive detection of fissile material,” *IEEE Trans. Nuclear Science*, 51, 2238-2244, 2004.
- [5] W.W. Black, “Application of correlation techniques to isolate structure in experimental data,” *Nuc. Inst. Meth.*, 71, 317, 1969.